



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 777 179 B1

B8

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention
of the grant of the patent:
15.05.2002 Bulletin 2002/20

(51) Int Cl.7: G06F 9/46

(21) Application number: 96307771.4

(22) Date of filing: 28.10.1996

(54) Direct bulk data transfers

Direkte Massendatenübertragung

Transfert direct de données en vrac

(84) Designated Contracting States:
DE FR GB IT SE

(30) Priority: 13.11.1995 US 556618

(43) Date of publication of application:
04.06.1997 Bulletin 1997/23

(73) Proprietor: Compaq Computer Corporation
Houston, Texas 77070-2698 (US)

(72) Inventors:

- Sprenkle, Todd W.
Sunnyvale, California 94086 (US)
- Murthy, Srinivasa D:
San Jose, California 95131 (US)
- Khatri, Anil
San Jose, California 95148 (US)

(74) Representative: Charig, Raymond Julian et al
Eric Potter Clarkson,
Park View House,
58 The Ropewalk
Nottingham NG1 5DD (GB)

(56) References cited:

- COMPUTER, vol. 27, no. 3, March 1994, LONG BEACH US, pages 75-82, XP000443073
KRIEGER ET AL: "The Alloc Stream Facility"
- 25TH INTERNATIONAL SYMPOSIUM ON FAULT-TOLERANT COMPUTING, 27 June 1995, PASADENA, USA, pages 2-11, XP000605839
BAKER ET AL: "A Flexible ServerNet-based Fault-Tolerant Architecture"
- IEEE MICRO, vol. 15, no. 1, February 1995, NEW YORK US, pages 37-45, XP000501487 HORST:
"TNet: A Reliable System Area Network"

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

[0001] The present invention is directed generally to direct data transfers in data processing systems, and more particularly to direct bulk input/output transfers done in a system via a network that provides connectivity for processor and input/output communication.

[0002] Third party transfer systems provide various schemes for data transfer. For example, High-Performance Storage System (HPSS) is an advanced, distributed Hierarchical Storage Management system that is capable of coordinating concurrent input/output (I/O) operations over a network to achieve high aggregate I/O throughput. HPSS is the next generation software system under development at the National Storage Laboratory. More information on HPSS can be obtained through the following World Wide Web Pages on the Internet:

http://www.ccs.ornl.gov/HPSS/HPSS_overview.html
http://www.llnl.gov/liv_comp/nsf/hpss/hpss.html

[15] This information discloses a Mover which is utilized for transferring data from a source device to a sink device in the HPSS. This Mover also performs a set of device control operation.

[0003] Fibre Channel (FC) also provides for third party transfer. FC allows for the transfer of data between workstations, mainframes, supercomputers, desktop computers, storage devices, displays and other peripherals. More information on FC can be obtained through the following World Wide Web Page on the Internet:

<http://www.amdahl.com/ext/CARP/FCA/FCA.html>

Finally, the IEEE Storage Systems Standards Working Group has created a draft Reference Model for Open Storage Systems Interconnection (OSSI). The following is from Version 5 of the draft OSSI document:

"2.3.3. Separation of Control and Data-flows

[0004] "The OSSI Model distinguishes control flows from data flows occurring between a client, a data source, and a data sink. Control flows carry requests, replies, and asynchronous notifications between a client and the data source or sink device. Control flows between the data source and the data sink carry source-sink protocol information to manage the flow of data. Data-flows pass only from a source to a sink. By logically separating control and data flows, the OSSI Model offers the possibility of optimizing each flow through separate implementation."

"2.3.4. Third-Party Transfers

[0005] "The OSSI Model allows data to flow directly between independent sources and sinks, under the control of a third party, initiating and controlling agent or client. Each entity separately performs operations such as data-flow control, error reporting, or initiating and terminating the transfer."

More information about OSSI can be obtained via the following World Wide Web Page:

<http://www.arl.mil/IEEE/ssswg.html>

[0006] U.S Patent No. 5,751,932 discloses a computing system which corresponds to the pre-characterising part of claim 1 and which responds to the need for a multiple processing system in a reliable system area network that provides connectivity for inter processor and input/output communication. The system of this patent is also reviewed in 25th International Symposium on Fault-Tolerant Computing, 27 June 1995, Pasadena, USA, pages 2-11, Baker et al: 'A Flexible Server Net-based Fault-Tolerant Architecture'.

This patent (referred to hereinafter as the referenced patent) teaches a system that provides a fail-fast, fail-functional, fault-tolerant microprocessor system and an architecture that includes a system area network (SAN) cloud formed by a number of router devices and associated interconnecting links, that enables any central processing unit (CPU) to communicate with any input/output (I/O) controller. This architecture is illustrated herein in Fig.1. Therefore, any I/O controller can be addressed by any CPU, and any CPU can be addressed by any I/O controller.

[0007] Broadly, the invention disclosed in the referenced patent includes a processing system composed of multiple sub-processing systems. Each sub-processing system has, as the main processing element, a CPU. This CPU comprises a pair of processors operating in a lock-step synchronized fashion to execute each instruction from an instruction stream simultaneously. Each of the sub-processing systems connects to an I/O system area network (SAN) that provides redundant communication paths between various components of the larger processing system. These various components include a CPUs and assorted peripheral devices (e.g., mass storage units, printers, and the like). These

redundant communication paths can also be between sub-processors that make up the larger overall processing system. Communication between any component of the overall processing system (e.g., between a first CPU and a second CPU, or between a CPU and any peripheral device) is implemented by forming and transmitting messages which are included in packets. In the preferred embodiment, each packet contains 64 bytes of data. These packets are routed from the transmitting or source component (e.g., a CPU) to a destination element (e.g., a peripheral device) by the SAN.

[0008] This system area network structure includes a number of router elements that are interconnected by a plurality of interconnecting links. The routers disclosed in the referenced patent which route packets from the source of the packet to the destination of the packet, do not themselves originate the packets. Routers act as packet switches by taking an incoming packet on one link and sending it out on the appropriate link for its destination. The router elements are responsible for choosing the proper or available communication paths from a transmitting component of the processing system to the destination component. The communication paths are based upon information contained in the message packet. Thus, the routing capability of the router elements provide the CPUs' I/O system with a communication path to the peripheral devices.

[0009] The architecture disclosed in the referenced patent uses a disk process pair to manage each disk, one half of the disk process pair will be the primary disk process and the other half will be the backup disk process. Additionally, the disk processes controlling disks on a SCSI chain are not confined to two CPUs, and the disk processes can be configured to run among multiple CPUs. When the SAN cloud of the referenced patent is being utilized, both CPUs and controllers can originate read and write cloud transactions for the CPU memory.

[0010] Fig. 1 illustrates an example of a storage architecture. This configuration shows CPUs 20, 22, 24 and 26 along with disk/storage controllers 30 and 32 are connected to SAN cloud 10. Controllers 30 and 32 include I/O packetizers 34 and 36 along with SCSI chips 40, 42, 44 and 46. I/O packetizers convert data packets from the network protocol into the bus protocol. This configuration also shows six SCSI disk pairs 100-105 and 110-115 hanging off two SCSI chains 120 and 122 (for a total of 12 storage disks 100-105 and 110-115). The disks are configured as primary disks 100 (\$A), 101 (\$B), 104 (\$E), 105 (\$F), 111 (\$D) and 114 (\$C) and mirror disks 102 (\$C), 103 (\$D), 110 (\$F), 112 (\$B), 113 (\$E) and 115 (\$A). The four CPUs (CPU0 20, CPU1 22, CPU2 24 and CPU3 26) accommodate disk processes which control the six disk pairs 100-105 and 110-115. The primary disk processes (\$A-P, \$B-P, \$C-P, \$D-P, \$E-P and \$F-P) and the backup disk processes (\$A-B, \$B-B, \$C-B, \$D-B, \$E-B and \$F-B) are scattered among the four CPUs 20, 22, 24 and 26. For example, the primary disk process 130 (\$A-P) for storage disk 100 (\$A) may be located in CPU0 20, and the backup disk process 133 (\$C-B) for storage disk 102 (\$C mirror) may be located in CPU1 22. Disk Processes 130-141 can be located (as shown in Fig. 1) in more than two CPUs. In another configuration, eight storage disk pairs can be hanging off two SCSI chains (for a total of 16 storage disks), and an extra SCSI chip can be located in each controller for external storage devices.

[0011] As an example of the operation of the system shown in Fig. 1 assume a request process located in CPU1 22 wishes to write to disk 100. To do so, the request process will first send a write data message to the (primary) disk process 130 located in CPU0 20 (disk process 130 controls disk 100). Disk process 130 then computes checksums over the data. Checksums ensure data integrity for the block of data being transferred. The block of data is then transferred from CPU0 20 to disk 100 through SAN cloud 10. When this transfer is completed, disk process 130 replies to request process 145.

[0012] Fig. 2 illustrates an example of disk transfers. Fig. 2 has most of the same elements included in Fig. 1. In addition to these elements, a buffer 150 is shown located in aCPU0 20 and a buffer 160 located in CPU1 22 are shown. In this example, the request process and the disk process are located in different CPUs. These processes can also be located in the same CPU. At step 1, the request process wants to write data from buffer 160 to \$A disk 100. Therefore, the data located in buffer 160 is sent to buffer 150 located in CPU0 20 because disk process 130, located within CPU0 20, is the disk process for \$A disk 100. At step 2, disk process 130 writes the data located in buffer 150 to disk 100. At step 3, disk process 130 replies to request process 145 because the transfer of data to disk 100 is complete.

[0013] In this arrangement, data is copied from the CPU with the request process to an intermediate CPU with the associated disk process and then from that intermediate CPU to the disk. During the transfer of data between CPUs and storage disks, it is desirable to remove similar data copies.

[0014] According to the present invention there is provided a data processing system for transferring data between a plurality of central processing units (CPUs) including a request CPU, at least one storage unit a network interconnecting said plurality of CPUs, one of said CPUs controlling access to said storage unit, said data processing system characterised by:

access means for providing direct access of said storage unit by said request CPU, said access means including means in said request CPU for creating a virtual memory address for a buffer memory of said request CPU and for providing said virtual memory address along with a storage unit access request to said one of said CPUs controlling access to said storage unit, means in said one of said CPUs controlling access to said storage unit for sending a work request including said virtual memory address to said storage unit, said storage unit responding

to said work request and interfacing directly with said request CPU through said network, wherein the data is transferred directly between said request CPU and said storage unit.

[0015] The invention will be further described by way of non-limitative example with reference to the accompanying drawings, in which:-

- Fig.1 illustrates an example of a storage architecture;
- Fig.2 illustrates an example of disk transfers;
- Fig.3 illustrates the transfer of data occurring when the controller is "pushing";
- Fig.4 illustrates an example of a direct data transfer for a disk read;
- Fig.5 illustrates an example of a direct data transfer for a disk write;
- Fig.6 illustrates how two node IDs are used for two paths to a disk and its mirror disk;
- Fig.7 is an example of system software layering with direct data transfers;
- Fig.8 illustrates the process flow with the read request path with direct data transfer;
- Fig.9 illustrates the process flow for the read reply path with direct data transfer;
- Fig. 10 illustrates the process flow for a write request path with direct data transfer; and
- Fig. 11 illustrates the process flow for a write reply path with direct data transfer.

[0016] The present invention provides for direct bulk data transfers in a reliable system via a network that provides connectivity for processor and I/O communication. This direct data transfer retains the disk process pair arrangement while eliminating the copying of data between the CPU running the request process and the CPU running the disk process. Thus, the data is copied directly between the request process' buffer and the storage unit. As a result, (1) the network bandwidth is saved because the data is not transferred to the buffer of the disk process, (2) the context switch time is saved because the disk process does not have to receive the data into its buffer, (3) the input/output (I/O) latency is reduced because a buffer copy is avoided, and (4) work is "off-loaded" from the message system because the message system does not send the data between the request process and the disk process. This message system is used for interprocessor communication.

[0017] Fig. 3 illustrates the transfer of data occurring when the controller 30 is "pushing." For example, when data is transferred from disk 105 to CPU memory, the controller 30 is "pushing" data 170 to CPU memory. This is also referred to as a "read" operation. In this illustration, both the request process and disk process are located in CPU3 26. This CPU3 26 sends controller 30 a virtual address 172 that identifies a CPU memory buffer 162 in CPU3 26 which will receive data 170, along with the parameters associated with the data transfer (e.g., a remote node identification). The virtual address 172 includes information used to obtain the physical (memory) address of buffer 162. This virtual address 172 will be later translated back into a physical address. In the preferred embodiment, the virtual address 172 includes both a page number and an offset which need to be used together to determine the associated physical address. Therefore, more than one virtual address 172 can be provided for one buffer. The virtual address is disclosed in greater detail in the referenced patent. Buffer 162 belongs to the request process. The parameters associated with the transfer are sent from CPU3 26 to controller 30 via a "work request." With this information, controller 30 can transfer data 170 directly from an I/O device (e.g., \$F disk 105) to the request process' memory. In a similar arrangement, the CPU can "write" to the I/O device (the disk). This is referred to as the controller "pulling" the data from the CPU memory. In this arrangement CPU3 26 sends adapter 30 the virtual address 172 for the CPU buffer 162 containing the data, along with the parameters associated with the I/O space. Again, these parameters are included in a "work request." With this information, controller 30 can transfer the data from CPU memory to disk device 105. Therefore, data is transferred between a CPU with a nondisk process and a disk without transferring the data through the CPU running the associated disk process.

[0018] Communications through the SAN cloud are by packets that contain a header with entries 174 and 176, respectively identifying the source node and destination node for the packet. This identification is in the form of a node identification (ID). In the preferred embodiment, only CPUs and other peripheral devices have node ID. Each packet header also contains a transaction type field (not shown in Fig. 3) that indicates the type of packet. When the packet type is a read or a write, the packet elicits an acknowledgement from the destination node. For read packets, this acknowledgement contains the data to be returned to the node which is performing the read. For the write packets, this acknowledgement returns with a successful or failure status for the write operation.

[0019] The virtual address 172 specified in the read/write packets is not a physical address at the destination node, but instead an address that is permission checked and then, if the permission check passes, is translated into a physical address. Permission checking is a form of, for example, source node identification validation, translation type checking (i.e., read/write permissions) and also bounds checking. An address validation and translation (AVT) table is used for the translation and permission checking applied to virtual addresses. In order to validate the source node ID, the source ID field of an accessed AVT entry specifies the source that corresponds to the AVT entry being used. This source ID

field is compared to the source ID contained in the requesting message packet, such that a mismatch will result in an AVT error interrupt which denies access. Transaction type checking involves checking to determine if access for either a read or a write can be granted. For bounds checking, the lower bound field and the upper bound field of the AVT entry is compared to an offset value to determine if access will be granted. This type of permission checking is described in greater detail in the referenced patent (e.g., see Figs. 13A-13C).

[0020] The hardware direct memory access (DMA) engine used to achieve the transfer of data is also referred to as the block transfer engine. In the preferred embodiment, the block transfer engine can compute checksums on memory buffers asynchronously, and at least one DMA engine is located in each CPU. As stated above, checksums ensure data integrity for the block of data being transferred. As stated above, any CPU 20, 22, 24 or 26 (Fig. 2) can delegate a buffer to the block transfer engine for transferring packets of data over the SAN cloud 10. In addition, any CPU 20, 22, 24 or 26 can delegate a buffer to the block transfer engine for computing the checksums and for depositing the checksums in another buffer. Finally, in the preferred embodiment, a DMA engine located in SCSI chip 40 performs the "pushing" or the "pulling" for each data transfer (for the disk "read" or "write").

[0021] As stated above, when the direct data transfer of the present invention takes place, the data is not transferred between the request originating CPU and the CPU containing the disk process. Instead, a virtual address is created and used. This virtual address is used by the controller to perform the direct data transfer. Again, there are two types of data transfers in this process. The first is a disk write operation which involves creating and using a virtual address. The controller uses this virtual address to do the necessary operations for the disk write. Thus, the AVT entry for the virtual address (for the request process' buffer) specifies the node ID for the controller which needs to access the request process' buffer.

[0022] For disk write operations, the originating CPU running the request process first computes the checksums of the data before issuing a request to the disk process. After the request process creates the virtual address for its buffer, the request process passes that virtual address to the disk process. The disk process cannot use this virtual address to access the request process' buffer because the source ID validation check would fail. After receiving the virtual address, the disk process sends a "work request" to the controller specifying (1) the handle to the node ID for the CPU originating the request (the CPU containing the request process) and (2) the virtual address for the buffer associated with the request process. The controller then "pulls" the data from the CPU memory buffer using the virtual address because this is a write operation. At the end of this transfer of data, a notification (typically in the form of an interrupt) is transmitted back to the CPU containing the disk process. The disk process then replies to the request process to notify it of the completion of the data transfer.

Additional information related to the AVT table and the virtual address is provided in the referenced patent.

[0023] The second type of data transfer is for a disk read. Fig. 4 illustrates an example of a direct data transfer for a disk read. At step 11, the virtual address for memory buffer 150 is created by the CPU 20 that will request the read. Buffer 150 is associated with the request process 146. At step 12, the request process 146 sends a request to disk (Dp2) process 133 resident on the processor CPU1 22. At step 13, disk process 133 sends a work request to controller 30. At step 14, controller 30 performs a transfer of data from \$F disk 105 to buffer 150. Step 14 is performed until all of the requested data has been sent to buffer 150. When the data transfer is complete, controller 30 interrupts disk process 133 at step 15. At step 16, disk process 133 informs request process 146 that the data transfer has ended. At step 17, the originating CPU running the request process computes the checksums of the buffer and then validates it before accepting the data.

[0024] Direct data transfers for disk writes are very similar to disk reads, except that (1) the checksums are computed by the request process before a request is sent to the disk process and (2) the direction of the transfer of data is from the CPU memory to the disk. Fig. 5 illustrates an example of a direct data transfer for a disk write. At step 21, the request process 146 performs the checksums calculations. At step 22, request processor 146 creates a virtual address for memory buffer 150 and for the checksums buffer. At step 23, request process 146 sends its request to disk process 133. At step 24, disk process 133 sends a work request to controller 30. At step 25, controller 30 transfers the data from buffer 150 and the checksums from the checksums buffer to \$F disk 105. After all of the requested data have been transferred from buffer 150 to disk 105, controller 30 interrupts disk process 133 at step 26. At step 27, disk process 133 informs request process 146 that the data transfer is complete.

[0025] The software for the direct data transfer allows the data to be sent directly to (or from) the CPU running the request process from (to) the disk device. To perform direct data transfer in the preferred embodiment, a device handle is utilized. This device handle provides the remote node ID to the AVT entry for the data transfer. Only the controller identified by this node ID can access the entry in the AVT table for the data transfer. The AVT table maps the requesting CPUs buffer to the virtual address space, so that it is accessible to the controller with the correct node ID. When activated, the device handle points to the data transfer parameters which include the node ID. To obtain this device handle, a code routine labeled "TSER_DEVINSTALL" is used in the preferred embodiment. Therefore, the parameters associated with the remote node are specified to the code when this routine is called.

[0026] In the preferred embodiment, the system initiates direct data transfers and performs a TSER_DEVINSTALL

when a request for data transfers is recognized. When the direct data transfer process is complete, the device handle is returned by calling, for example, the TSER_DEVREMOVE routine. In a second embodiment, either global device handles or stack/module/slot to device handle translation caches are located in each CPU. In this embodiment, the TSER_DEVINSTALL and TSER_DEVREMOVE routines are not invoked.

5 [0027] Users can indicate their desire to use direct data transfers in several ways. For example, SETMODE 141 can enable a large data transfer mode for disk file transfers. This SETMODE enables users to specify that they desire large, unstructured accesses to disk files that bypass the disk process cache. In another example, BULKREAD and BULKWRITE can be used as internal tools when a backup, restore, dump, or the like is requested. BULKREAD and BULKWRITE also bypass the disk process cache. Thus, if a BULKREAD or a BULKWRITE or a SETMODE 141 is enabled for file transfers, direct data transfer will be used to perform the transfer task.

10 [0028] In the preferred embodiment, when SETMODE 141 is indicated, a message is sent to the disk process to flush its cache. As a part of the reply to this SETMODE request, the disk process returns an indication to the file system that it can support the direct data transfer. The disk process also returns two IDs to the file system. These IDs constitute the two primary paths to the relevant disk and its mirror disk. Both IDs are used when doing a write to the disk devices, 15 and only one ID is used for a read operation. When the file system receives the response from the disk process, it deciphers the reply and calls a direct data transfer routine to install the two IDs with the code.

15 [0029] Similarly, when BULKREAD or BULKWRITE is indicated, the file system sends a normal bulk data transfer request to the disk process. When possible, the disk process responds to the file system indicating that it can support the direct data transfer. At this time, the disk process also returns the two IDs that constitute the two primary paths to 20 the relevant disk and its mirror. The file system then uses direct data transfer for subsequent bulk data transfer calls, and also calls a direct data transfer routine to install the two IDs with the server net code.

25 [0030] As part of the direct data transfer session establishment (installing the two IDs), the direct data transfer routine also acquire space for a transfer information block (TIB) from the FLEXPOOL. The TIB is used for the interface between the direct data transfer routines and the code. The TIB is also used by the direct data transfer routines to request that the code use the DMA engine to do the checksums calculations. In the preferred embodiment, a single TIB is used because the checksums operation is synchronous. The FLEXPOOL allocates space for memory management.

30 [0031] Fig. 6 illustrates how two IDs are used for two paths to a disk and its mirror disk. In the preferred embodiment, the two IDs are associated with SCSI controller chips 40 and 42. As shown in Fig. 6, these SCSI controller chips 40 and 42 provide access to disk 100 and its mirror disk 115 respectively. In the preferred embodiment, a fault tolerant system is provided with the following: CPU0 20 and CPU1 22; two SAN clouds 10 and 12; a primary path 14 to disk and its mirror and a backup path 16 to disk and its mirror controllers 30 and 32; SCSI controller chips 40, 42, 44 and 46; a disk 100 and its mirror disk 115; and chains 120 and 122. In this arrangement, all components of the system have at least one redundant component as a backup (e.g., disk 100 has a backup mirror disk 115). Therefore, during 35 a data write operation, all data transfers are made to two disks (a disk and its mirror) such that if an error occurs and correct data cannot be accessed from one disk, that data will also be located on another disk for access. This configuration results in very little, if any, data corruption. During a disk read operation, direct data transfer software maps the data buffer for access by only one of the disks 100 or 115. Thus, the buffer will be mapped into only one virtual address space.

40 [0032] In an alternative embodiment, four server net IDs are used for direct data transfers. In this arrangement, data buffers must be redundantly mapped into the virtual address space for access by all four SCSI controller chips 40, 42, 44 and 46. This enables the disk process to issue the data transfer request to any of the SCSI controller chips 40, 42, 44 or 46.

45 [0033] The direct data transfer support software is structured as a library of software routines (referred to as routines). In the preferred embodiment, the file system is a client to these routines. Thus, when the file system determines that a direct data transfer is being initiated, it calls the routines which take necessary actions to map buffers and do checksums calculations. When the direct data transfer routines are done, the file system uses the message system to send a message to the appropriate CPU containing the disk process.

50 [0034] Fig. 7 is an example of system software layering with direct data transfers. As illustrated in Fig. 7, the file system 210 can call the message system 230 directly and can call the routines in direct data transfer library 220 before invoking the message system. The file system 210 calls the direct data transfer library 220 when the flag enabling direct data transfer is set. As mentioned earlier, this flag becomes set when a SETMODE 141 or a BULKREAD/BULK-WRITE is performed.

55 [0035] If the direct data transfer is enabled at FILE_CLOSE time the file system calls a direct data transfer routine to end the direct data transfer session. This allows the direct data transfer routines to deinstall the two IDs with the code. FILE_CLOSE time indicates no more files will be transferred.

[0036] In another embodiment of the present invention, a device handle caching scheme is used when multiple processes in the same CPU are doing direct data transfers to the same disk or tape. This avoids the replication of device handles that point to the same node ID in the CPU. When a device handle caching system is used, the direct

data transfer routine checks to see if a device handle exists in the CPU when the direct data transfer session is established. If a device handle already exists, that device handle is used. Otherwise, a new device handle is created by calling the code. At FILE_CLOSE time, the device handle is not removed.

[0037] For read data transfer operations, the checksums and verification is done after completion of the data transfer.

- 5 In the preferred embodiment, the direct data transfer initiation routine (1) allocates a checksums buffer, (2) maps the request processor's buffer into virtual address space, and (3) returns to the file system. The mapping is done to enable the controller to push the data over the primary path in the SAN. This mapping operation involves a processor cache sweep which makes the memory is in the request CPU consistent with the processor's cache. The file system embeds the mapped virtual addresses in its request control area. The related message is then sent to the appropriate media server process (disk process or tape process). This message from the request process indicates that the transfer is a direct data transfer. The disk process then knows that the reply data buffer and checksums buffer are mapped for access by the controller. The disk process issues a command to the controller telling it (1) what operation to perform and (2) to deposit the read data and checksums in the destination CPU (CPU containing the request process) at the specified addresses. The controller interrupts the CPU containing the disk process when the data transfer is complete.

10 15 This arrangement was illustrated in Fig. 4.

[0038] In the preferred embodiment, the mapping of the request processor's buffer into virtual address space is done by a direct data transfer routine. These virtual addresses are transferred to the disk process. To facilitate this transfer, the direct data transfer routine returns the virtual addresses to the file system which embeds the virtual addresses in the request it sends to the CPU containing the disk process. When the data transfer is complete, the controller notifies the disk process via an interrupt that the request is done. When the request process receives the message that the data transfer is complete, it calls a direct data transfer routine to handle the completion of the read request. The called direct data transfer routine then calls the code to calculate the appropriate checksums for the request CPUs buffer. This checksums operation is a blocking operation. Therefore, the process is suspended while the checksums operation is in progress. The routine unblocks the calling process when the checksums calculations are complete. The direct data transfer routine then verifies that the calculated checksums match the checksums returned by the controller. The routine then returns either a success or checksums error indication to the calling file system routine.

[0039] Fig. 8 illustrates the process flow with the read request path with direct data transfer. At step 300, the application calls the file system for a read request. At step 302, the file system prepares a request section and calls the direct data transfer routine. At step 304, the direct read initiation routine allocates a checksums buffer from FLEXPOOL.

- 20 At step 306, the direct read initiation routine maps the request CPUs buffer to enable access of the buffer to the controller over the SAN and returns the virtual addresses to the file system. At step 308, the file system calls the message system to send a message to the disk process. This message includes the buffers' virtual addresses which are embedded in the request section. At step 310, the disk process issues a work request to the controller. This work request is for depositing the read data and checksums in the CPU containing the request process. At step 312, the controller performs the data transfer and then interrupts the disk process with the completion notification. At step 314, the disk process replies to the request process' message.

[0040] Fig. 9 illustrates the process flow for the read reply path with direct data transfer. At step 320, the disk process' reply causes the request process' message to be queued, and the request process is woken up. At step 322, the request process wakes up and calls the direct read completion routine to verify the checksums. At step 324, the direct read completion routine calls the code to queue the checksums calculations for the DMA engine. The code then suspends the calling process. The DMA Engine then performs the checksums calculations. At step 326, the DMA engine checksums calculations completion interrupt causes the code to unblock the request process. At step 328, the direct data transfer routine verifies the calculated checksums with the checksums deposited by the controller. At step 330, the direct data transfer routine returns the checksums buffer to FLEXPOOL and returns with the success or failure of the checksums comparisons to the file system. At step 332, in the preferred embodiment, the file system returns to the request process if the comparison is successful, or retries without direct data transfer if a checksums error occurs.

[0041] For write data transfer operations, the checksums are calculated before the data is transferred. The file system calls the direct write initiation routine and provides the address of the request CPU's buffer. The direct data transfer routine first allocates a checksums buffer and maps the relevant buffers (the request data and checksums buffers) into a virtual address space. For the write operation, the buffers are mapped into the virtual address space for access by two paths of a disk and its mirror. If this is a mirrored write, both of the related controllers will access these buffers. A mirrored write is used in the preferred embodiment.

[0042] The direct data transfer routine then calls the code to calculate the checksums of the buffer. The code then queues the checksums calculations to the request CPU's DMA engine and blocks the calling process until completion of the checksums calculations. The DMA engine computes the specified checksums and deposits the resultant checksums in the checksums buffer. A completion interrupt for the checksums calculations is then fielded by the code. The completion interrupt unblocks the request process at that time.

[0043] The direct data transfer routine resumes processing by returning four virtual addresses to the file system.

Two of these virtual addresses corresponds to the request CPU's buffer. The other two virtual addresses correspond to the checksums buffer. The file system write routine embeds these four virtual addresses into the request control section of its message and then calls the message system to send the associated message to the disk process. In the preferred embodiment, in order to avoid data corruption, the request process buffers are mapped to only one controller when a disk read is performed.

[0044] The disk process, upon receipt of this direct data transfer write request, sends the appropriate work request to the SCSI controllers. This work request specifies that the data and checksums buffers come from the CPU containing the request process, but that the request completion interrupt from the controller should go to the CPU containing the disk process. The controller then pulls the data for the write operation from the request process' CPU. The data is then written to the physical medium, and the disk process is notified via an interrupt upon completion. The disk process waits for the completion interrupt from both the primary and mirror disk halves. After both disks have completed the operation, the disk process replies to the request process' message.

[0045] When the reply sent by the disk process arrives at the request process's CPU, the CPU wakes up the request process. The file system then calls the direct data transfer routine to deal with the write completion. In this case, the direct data transfer routine unmaps the buffers from the virtual address space, deallocates the checksums buffer and returns to the file system. The file system then notifies its user of the completion of the write request.

[0046] Fig. 10 illustrates the process flow for a write request path with direct data transfer. At step 350, the application calls the file system for the write operation. At step 352, the file system calls the direct data transfer write initiate routine. At step 354, the direct data transfer write initiate routine allocates a checksums buffer from the FLEXPOOL. At step 356, the direct data transfer write initiation routine maps the request CPU's buffer and the checksums buffer for access by the two controller paths. The routine then calls the code to do the checksums calculations. At step 358, the code queues the checksums calculations to the DMA engine and suspends the request process. At step 360, the completion interrupt causes the code to resume the request process and to return to the direct write initiation routine. At step 362, the direct data transfer write initiation returns the virtual addresses of the request and checksums buffers to the file system. At step 364, the file system embeds these virtual addresses in its request control area and sends this message to the disk process.

[0047] Fig. 11 illustrates the process flow for a write reply path with direct data transfer. At step 370, the code wakes up the request process. At step 372, the file system calls the direct data transfer write completion routine. At step 374, this direct data transfer routine unmaps the buffers from the virtual address space, deallocates the checksums buffer and returns. At step 376, the file system returns a write completion notification to the application.

[0048] In the preferred embodiment, the space for the checksums buffers is allocated by the direct data transfer library routines. This space is allocated from the FLEXPOOL. In an alternative embodiment, the file system allocates the space for the checksums buffers. The size of the checksums buffer depends on the request CPU's buffer size and the sector/block size.

[0049] In the preferred embodiment, the file system client who initiated the direct data transfer operation can also cancel that operation. This cancellation can be done at the end of a data transfer or while outstanding data transfer requests are present. When a direct data transfer operation is cancelled in this manner, the file system calls a direct data transfer routine to perform the cancellation. This routine is responsible for unmapping the buffers from the virtual address space and for returning the associated checksums buffer to their pool. The file system then sends a cancel notification to the appropriate disk process. After these two operations are complete, the direct data transfer operation is deemed cancelled. When the direct data transfer routine unmaps a buffer, the controller will encounter errors if it tries to access this buffer. These errors are reported to both the controller and the CPU containing the request process. In the preferred embodiment, the controller stops accessing the buffer after the first error is encountered.

[0050] In the preferred embodiment, the node ID of all the CPUs in the system are registered with the firmware. This registration is in the form of a "set parameters" mailbox command in the firmware. In response to this mailbox command, the firmware returns an 8-bit handle. This handle is used by the SCSI module driver to identify the host processor to the firmware. This registration of node IDs allows for direct data transfer between request processors (processors containing the request process) and storage units (units containing disks).

[0051] The examples discussed thus far only consider the case when the request process and the disk process are located in different CPUs. If these two processes are located in the same CPU, direct data transfer can still be utilized. When direct data transfer is used in this example, the CPU copy of the buffer between the request process and the disk process is still avoided. Therefore, direct data transfer is also desirable when both the request process and disk process are located in the same CPU.

[0052] While disks were used as the storage devices in the above examples, any storage device can be used for the direct data transfer. For example, direct data transfer can be used on tape I/O to avoid unnecessary data copies. If that I/O does not include the checksums calculations, direct data transfer would be simpler than in the disk case. If the tape process also supports direct data transfer, the backup and restore utilities can take advantage of this feature. For example, when the disk process, the tape process and the backup/restore process are all scattered in three different

CPU, direct data transfers would avoid two unnecessary buffer copies.

[0053] In the preferred embodiment, the following software code excerpts provide the interface routines in the direct data transfer library. The first excerpt is for the direct data transfer session start routine (see directly below this paragraph). This routine is called by the file system when it receives an indication from a disk process that it can support direct data transfer requests. The disk process responds to the SETMODE 141/bulk data transfer requests with this indication. The response from the disk process also contains information used by the direct data transfer routines, such as node IDs, packetizer types, and the like. This information is passed on to the direct data transfer session start routine by the file system.

```

10
  /*
   * Returns 0 upon success
   * an error code upon failure
   */
15   iop_cookie is an input pointer to the area that the IOP returned to the
   * File System to pass on to the Direct IO routines.
   * The File System should not care
   * about the internal format of this area. The format of this area would be
   * an agreement between the IOPs and the Direct IO routines. This area
20   * will contain the Tnet IDs of the controller, the type of the packetizer
   * in the controller, etc.
   *
   * directio_cookie is a pointer to an area where directio can store some
25   * of its context for the directio session. Aside from allocating storage for
   * this area, the File System should not care about the format of this area.
   *
   * The size of the cookie areas is TBD.
   */
30   int DirectIO_Session_Start(void *iop_cookie, void *directio_cookie);
{
35     look at iop cookie and see how many tnet ids are in it.
     call tser_dev_install to install these tnet ids. Ask tser_dev_install not to
     worry about allocating interrupt and barrier AVTs.
     call tser_dev_set_packetizer to set the packetizer to the type specified
     in the iop_cookie.
     call tser_dev_set_tnetid to set the tnet id in the installed device handle.
40     store the returned device handles in the directio cookie area.
     allocate a Tib from FLEXPOOL and store the Tib address in the
     directio_cookie area.
}

45 [0054] The direct data transfer session end routine is called by the file system at FILE_CLOSE time or when a
      SETMODE 141 is done to disable large data transfers.

```

50

55

```

5      " Returns 0 upon success
      " an error code upon failure.
      "
10     " directio_cookie is a pointer to the area that was initialized during the
      " directio_session_start call.
      "
15     int DirectIO_Session_End(void *directio_cookie)
      {
          call tser_dev_remove to uninstall the Tnet IDs.
          deallocate the TIB buffer by returning it to the FLEXPOOL.
      }

```

[0055] The direct read start routine is called by the file system prior to initiating a direct data transfer read request.

```

20
25      "
      " Returns 0 upon success
      " an error code upon failure.
      "
30      " directio_cookie is a pointer to the direct io session information that was
      " initialized during directio session start.
      " buffer is a pointer to the buffer into which data should be read.
      " buffer_size is the size of the above buffer.
35      " tnet_vaddrs is a pointer to an array of two 32 bit integers. Two tnet
      " virtual addresses will be deposited there by Direct_Read_Start.
      " These two tnet vaddrs then need to be copied to the Request Control
      " and sent on to the IOP.
      "
40      " xsum_buffer is a pointer to a 32 bit integer location. Direct_Read_Start
      " deposits the address of the xsum buffer it allocated in this location.
      " This address needs to be passed to Direct_Read_End or
      " Direct_Read_Abort.
      "
45      int Direct_Read_Start(void *directio_cookie, void *buffer, int buffer_size,
                           void *tnet_vaddrs, void *xsum_buffer)
      {
          "xsum_buffer = allocate the checksum buffer from the Flex Pool.
          map the user buffer and the checksum buffer for write access by the
          device handle stored in the directio_cookie.
          deposit the two tnet virtual addresses returned by the map routine into the
          area pointed at by tnet_vaddrs.
50      }

```

[0056] The direct read end routine is called by the file system when it receives a response to the direct data transfer read request from the disk process. If the disk process response indicates success, the direct read end routine will do the checksums verification. In all cases, the direct read end routine will unmap the buffers and deallocate the checksums buffer.

```

    "
    " Returns 0 upon success
    " an error code upon failure. One such failure could be checksum mismatch.
5
    "
    " tnet_vaddrs is a pointer to the two tnet virtual addresses that were created
    " during the call to Direct_read_start.
    " iop_returned_status is the status code that the IOP returned in its response to
10
    " the read request sent by the File System.
    " xsum_buffer is the 32 bit address of the checksum buffer that Direct_Read_Start
    " allocated.
    "
15 int Direct_Read_End(void *directio_cookie, void *tnet_vaddrs, int iop_returned_status,
    void xsum_buffer)
{
    "
    unmap the two tnet virtual addresses listed in the tnet_vaddrs array.

20 if (iop_returned_status == success)
{
    "
    use the TIB from the directio cookie area.
    call tser_transfer to calculate checksum.
    compare device returned checksum to calculated checksum.
25
    set return value.
}
return xsum_buffer to FLEXPOOL.
return return_value
}
30

```

[0057] The direct write start routine is called by the file system prior to initiating a direct data transfer write request.

35

40

45

50

55

```

1
2      " Returns 0 upon success.
3      " otherwise an error code is returned.
4
5      " directio_cookie is a pointer to the area of memory that belongs to Direct IO
6      " that is allocated by the File System.
7
8      " buffer is a pointer to the user buffer that is being written to the media.
9      " buffer_size is the size of the above buffer.
10     " tnet_vaddrs is a pointer to an array of four 32 bit integers. The array needs to
11     " be allocated by the file system. Direct IO routines will deposit four Tnet
12     " virtual addresses into this array - two for the xsum buffer and two for the user
13     " buffer. The FS needs to copy these four into the request control it sends to
14     " the IOP.
15     " xsum_buffer is a pointer to a 32 bit location where Direct_Write_Start will
16     " deposit the address of the allocated xsum buffer. This address needs
17     " to be passed to Direct_Write_End or Direct_Write_Abort.
18
19 }
20 int Direct_Write_Start(void *directio_cookie, void *buffer, int buffer_size,
21                      void *tnet_vaddrs, void *xsum_buffer)
22 {
23     " xsum_buffer = allocate a xsum buffer from the FLEXPOOL.
24     " call the Tnet millicode map routines to map the xsum and user buffers for
25     " access by the disk and its mirror.
26     " deposit the Tnet virtual addresses returned above into tnet_vaddrs array.
27     " use the TIB whose address we stored in the directio_cookie.
28     " call tser_transfer to calculate checksum of buffer.
29
30 }

```

35 [0058] The direct write end routine is called by the file system when a reply to a direct data transfer write request is received from a disk process.

```

40     " Returns 0 upon success.
41     " an error code is returned otherwise.
42
43     " directio_cookie is a pointer to the area of File System memory where directio
44     " session information is stored.
45     " tnet_vaddrs is a pointer to an array of four 32 bit ints. This is the same as the
46     " array which was passed into direct_write_start.
47
48 int Direct_Write_End(void *directio_cookie, void *tnet_vaddrs, void *xsum_buffer)
49 {
50     " call tnet millicode to unmap the four tnet virtual addresses listed in
51     " tnet_vaddrs.
52     " deallocate xsum_buffer.
53
54 }

```

[0059] The direct read abort routine is called by the file system after it has cancelled the direct data transfer read request which it has sent to the disk process. The disk process is informed of the cancellation, and then the direct data

transfer routine is called to clean up the request.

```

5   /*
6    * Returns 0 upon success.
7    * an error code is returned otherwise.
8    * directio_cookie is a pointer to the area of the File System memory where directio
9    * session information is stored.
10   * tnet_vaddrs is a pointer to an array of two 32 bit ints. This is the same as the
11   * array which was passed into direct_read_start.
12   * xsum_buffer is the 32 bit address of the checksum buffer that Direct_Read_Start
13   * allocated.
14   */
15  int Direct_Read_Abort(void *directio_cookie, void *tnet_vaddrs, void *xsum_buffer)
16  {
17      call tnet millicode to unmap the two tnet virtual addresses listed in
18          tnet_vaddrs.
19      deallocate xsum_buffer.
20  }

```

25 [0060] The direct write abort routine is called by the file system after it has cancelled the direct data transfer write message which it sent to the disk process. The disk process is informed about the cancellation, and then the direct data transfer routine is called to clean up the request.

```

30  /*
31   * Returns 0 upon success.
32   * an error code is returned otherwise.
33   *
34   * directio_cookie is a pointer to the area of File System memory where directio
35   * session information is stored.
36   * tnet_vaddrs is a pointer to an array of four 32 bit ints. This is the same as the
37   * array which was passed into direct_write_start.
38   * xsum_buffer is the 32 bit address of the checksum buffer allocated by
39   * Direct_Write_Start.
40   */
41  int Direct_Write_Abort(void *directio_cookie, void *tnet_vaddrs, void *xsum_buffer)
42  {
43      call tnet millicode to unmap the four tnet virtual addresses listed in
44          tnet_vaddrs.
45      deallocate xsum_buffer.
46  }

```

50 [0061] As stated above, the checksums are calculated using the block transfer engine within each CPU. Thus, the checksums can be calculated asynchronously, without CPU intervention. The software sub-system in every CPU is responsible for performing checksums calculations. In an alternative embodiment, the checksums calculations could be performed in the file system code. In this example, an interrupt handler is used to handle the completion interrupt of the checksums calculations because the code only provides an asynchronous interface. In yet another embodiment, the code provides a synchronous interface. In this arrangement, the checksums can be performed in the file system layer without any interrupt handler.

[0062] While a full and complete disclosure of the invention has been made, it will become apparent to those skilled

In this art that various alternatives and modifications can be made to various aspects of the invention without departing from the scope of the claims which follow.'

5 **Claims**

1. A data processing system for transferring data between a plurality of central processing units CPUs (20,22), including a request CPU (20), at least one storage unit (30,34,105), and a network (10) interconnecting said plurality of CPUs, one of said CPUs (22) controlling access to said storage unit, said data processing system characterised by:

10 access means (220) for providing direct access of said storage unit by said request CPU (20), said access means including means (146) in said request CPU for creating a virtual memory address for a buffer memory (150) of said request CPU and for providing said virtual memory address along with a storage unit access request to said one of said CPUs controlling access to said storage unit (20), means (133) in said one of said CPUs controlling access to said storage unit for sending a work request including said virtual memory address to said storage unit, said storage unit responding to said work request and interfacing directly with said request CPU through said network, wherein the data is transferred directly between said request CPU and said storage unit.

- 20 2. The data processing system of claim 1, wherein said data access allows said request CPU (20) to read data from said storage unit (30,34,105) into said buffer memory (150) at said virtual memory address.

- 25 3. The data processing system of claim 1, wherein said data are transmitted in data packets, each of said data packets including a destination node identification, a source node identification, said virtual memory address, and a plurality of data words.

- 30 4. The data processing system of claim 3, wherein said storage unit (30,34,105) includes means (30) for advising said one of said CPUs controlling access to said storage unit (22) after each of said data packets containing the data for transfer is transmitted to said request CPU (20).

- 35 5. The data processing system of claim 1, wherein said direct access allows said request CPU (20) to write into said storage unit (30,34,105), said storage unit accessing said memory buffer (150) at said virtual memory address for transfer of said data.

- 40 6. The data processing system of claim 5, wherein said storage unit includes two disks for storing the same data, said direct access request including addresses to said two disks whereby said data is written into both of said two disks.

- 45 7. The data processing system of claim 1, wherein said access means includes at least one software routine for creating said virtual memory address for said buffer of said request CPU and for providing said virtual memory address along with said storage unit access request to said one of said CPUs.

45 **Patentansprüche**

1. Datenverarbeitungssystem zum Übertragen von Daten zwischen mehreren Zentraleinheiten (Central Processing Units / CPUs) (20, 22), mit einer Anforderungs-CPU (20), mindestens einer Speichereinheit (30, 34, 105) und einem die mehreren CPUs verbindenden Netzwerk (10), wobei eine der CPUs (22) den Zugriff auf die Speichereinheit steuert, wobei das Datenverarbeitungssystem gekennzeichnet ist durch:

50 eine Zugriffseinrichtung (220) zum Vorsehen eines direkten Zugriffs der Anforderungs-CPU (20) auf die Speichereinheit, wobei die Zugriffseinrichtung aufweist: eine Einrichtung (146) in der Anforderungs-CPU zum Erzeugen einer virtuellen Speicheradresse für einen Pufferspeicher (150) der Anforderungs-CPU und zum Liefern der virtuellen Speicheradresse zusammen mit einer Speichereinheit-Zugriffsanforderung an die eine den Zugriff auf die Speichereinheit (20) steuernde CPU, eine Einrichtung (133) in der einen den Zugriff auf die Speichereinheit steuernden CPU zum Senden einer die virtuelle Speicheradresse enthaltenden Arbeitsanforderung an die Speichereinheit, wobei die Speichereinheit auf die Arbeitsanforderung reagiert und über das

Netzwerk mit der Anforderungs-CPU direkt Verbindung aufnimmt, wobei die Daten direkt zwischen der Anforderungs-CPU und der Speichereinheit übertragen werden.

2. Datenverarbeitungssystem nach Anspruch 1, bei dem der Datenzugriff der Anforderungs-CPU (20) erlaubt, Daten von der Speichereinheit (30, 34, 105) an der virtuellen Speicheradresse in den Pufferspeicher (150) zu lesen.
3. Datenverarbeitungssystem nach Anspruch 1, bei dem die Daten in Datenpaketen übertragen werden, wobei jedes der Datenpakete eine Zielknotenidentifikation, eine Ausgangsknotenidentifikation, die virtuelle Speicheradresse und mehrere Datenwörter enthält.
4. Datenverarbeitungssystem nach Anspruch 3, bei dem die Speichereinheit (30, 34, 105) eine Einrichtung (30) aufweist, die es der den Zugriff auf die Speichereinheit steuernden CPU (22) mitteilt, wenn jedes der die zu übertragende Daten enthaltenden Datenpakete an die Anforderungs-CPU (20) übertragen wurde.
5. Datenverarbeitungssystem nach Anspruch 1, bei dem der direkte Zugriff es der Anforderungs-CPU (20) erlaubt, in die Speichereinheit (30, 34, 105) zu schreiben, wobei die Speichereinheit zum Übertragen der Daten auf den Speicherpuffer (150) an der virtuellen Speicheradresse zugreift.
6. Datenverarbeitungssystem nach Anspruch 5, bei dem die Speichereinheit zwei Platten zum Speichern der gleichen Daten aufweist, wobei die direkte Zugriffsanforderung Adressen an die beiden Platten enthält, wodurch die Daten auf die beiden Platten geschrieben werden.
7. Datenverarbeitungssystem nach Anspruch 1, bei dem die Zugriffseinrichtung mindestens eine Software-Routine zum Erzeugen der virtuellen Speicheradresse für den Puffer der Anforderungs-CPU und zum Liefern der virtuellen Speicheradresse zusammen mit der Speichereinheit-Zugriffsanforderung an die eine der CPUs aufweist.

Revendications

1. Système de traitement de données destiné à transférer des données entre une pluralité d'unités centrales de traitement CPU (20, 22), incluant une CPU de requête (20), au moins une unité de mémoire (30, 34, 105), et un réseau (10) interconnectant ladite pluralité de CPU, une desdites CPU (22) commandant l'accès à ladite unité de mémoire, ledit système de traitement de données étant caractérisé par :
 - des moyens d'accès (220) destinés à fournir, à ladite CPU de requête (20), un accès direct à ladite unité de mémoire, lesdits moyens d'accès incluant des moyens (146) dans ladite CPU de requête pour créer une adresse de mémoire virtuelle relative à une mémoire tampon (150) de ladite CPU de requête et pour transmettre ladite adresse de mémoire virtuelle conjointement avec une requête d'accès à l'unité de mémoire à ladite une desdites CPU commandant l'accès à ladite unité de mémoire (30), des moyens (133) dans ladite une desdites CPU commandant l'accès à ladite unité de mémoire pour émettre une requête de tâche incluant ladite adresse de mémoire virtuelle à destination de ladite unité de mémoire, ladite unité de mémoire répondant à ladite requête de tâche et interagissant directement avec ladite CPU de requête au travers dudit réseau, dans lequel les données sont transférées directement entre ladite CPU de requête et ladite unité de mémoire.
 - Système de traitement de données selon la revendication 1, dans lequel ledit accès aux données permet à ladite CPU de requête (20) de lire les données de ladite unité de mémoire (30, 34, 105) à ladite mémoire tampon (150), à ladite adresse de mémoire virtuelle.
 - Système de traitement de données selon la revendication 1, dans lequel lesdites données sont transmises en paquets de données, chacun desdits paquets de données incluant une identification de noeud de destination, une identification de noeud de source, ladite adresse de mémoire virtuelle et une pluralité de mots de données.
 - Système de traitement de données selon la revendication 3, dans lequel ladite unité de mémoire (30, 34, 105) inclut des moyens (30) destinés à avertir ladite une desdites CPU commandant l'accès à ladite unité de mémoire (22) après que chacun desdits paquets de données contenant les données à transférer ait été transmis à ladite CPU de requête (20).
 - Système de traitement de données selon la revendication 1, dans lequel ledit accès direct permet à ladite CPU

de requête (20) d'écrire dans ladite unité de mémoire (30, 34, 105), ladite unité de mémoire accédant à ladite mémoire tampon (150), à ladite adresse de mémoire virtuelle, pour transférer lesdites données.

- 5 6. Système de traitement de données selon la revendication 5, dans lequel ladite unité de mémoire inclut deux mémoires à disque destinées à mémoriser ces mêmes données, ladite requête d'accès direct incluant des adresses auxdites deux mémoires à disque au moyen desquelles lesdites données sont écrites dans les deux dites mémoires à disque.
- 10 7. Système de traitement de données selon la revendication 1, dans lequel lesdits moyens d'accès incluent au moins un sous-programme logiciel pour créer ladite adresse de mémoire virtuelle relative à ladite mémoire tampon de ladite CPU de requête, et pour transmettre ladite adresse de mémoire virtuelle conjointement avec ladite requête d'accès à l'unité de mémoire à ladite une desdites CPU.

15

20

25

30

35

40

45

50

55

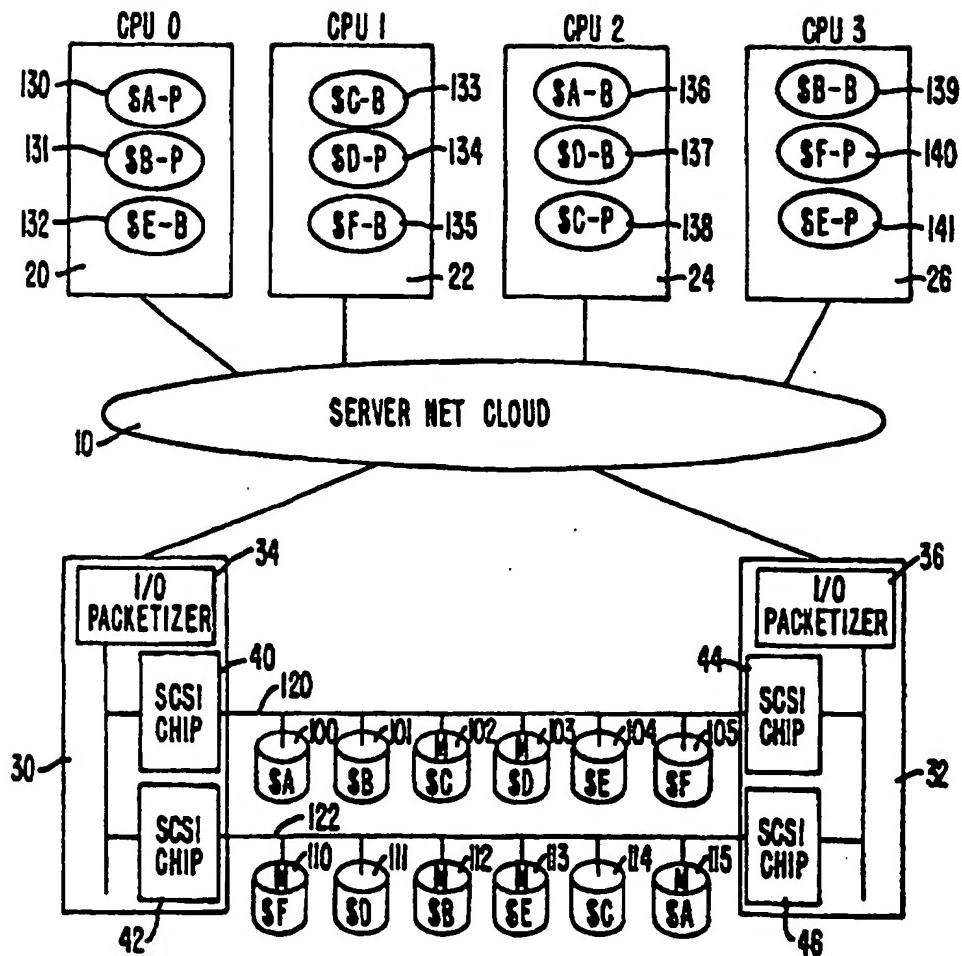


FIG. 1.

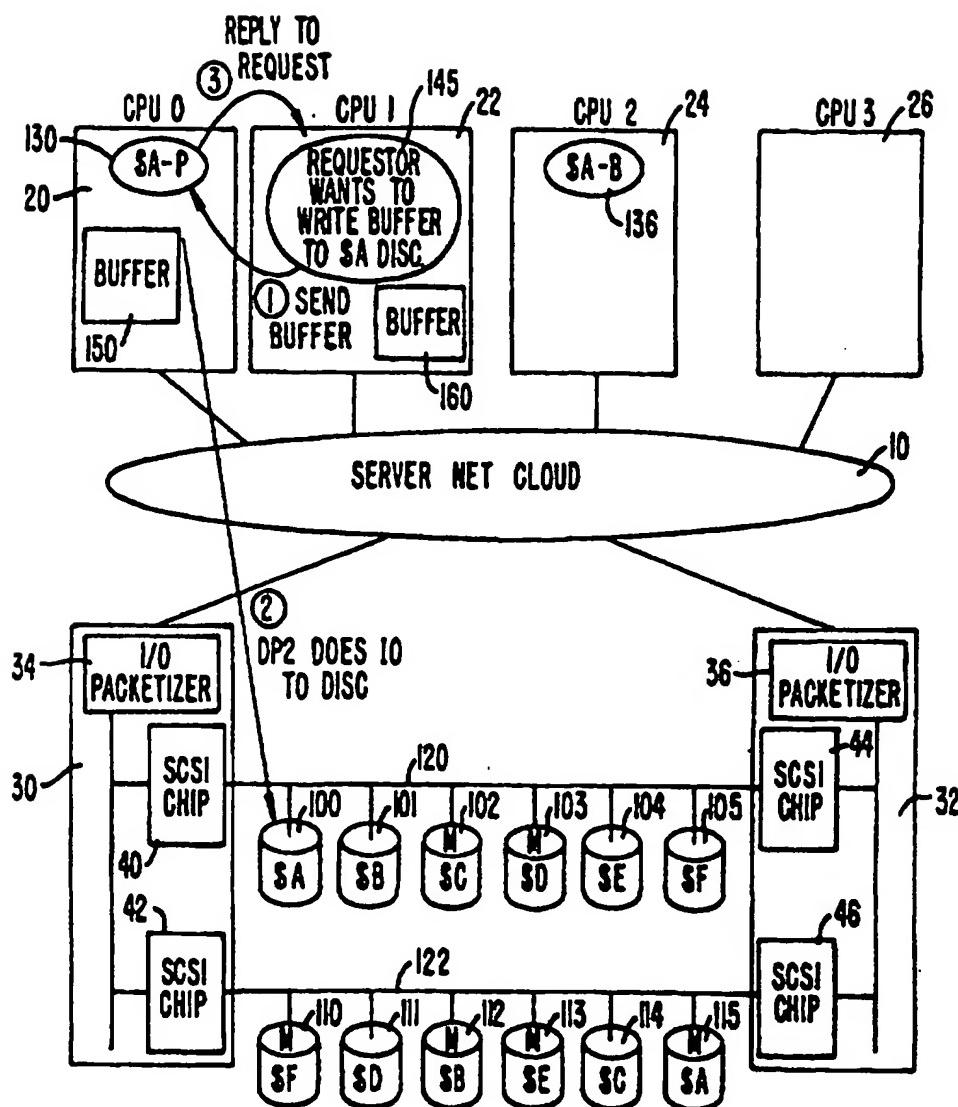


FIG. 2.

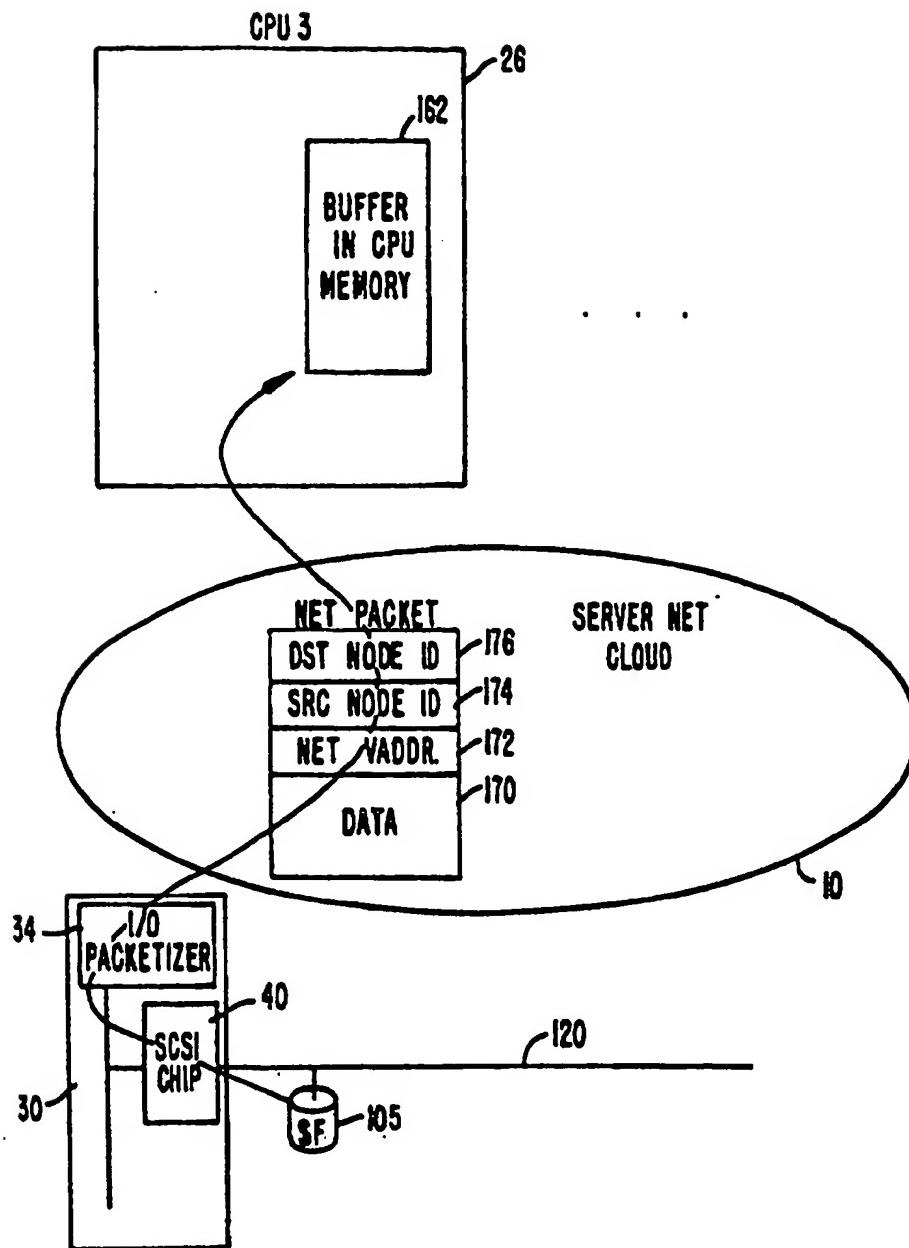
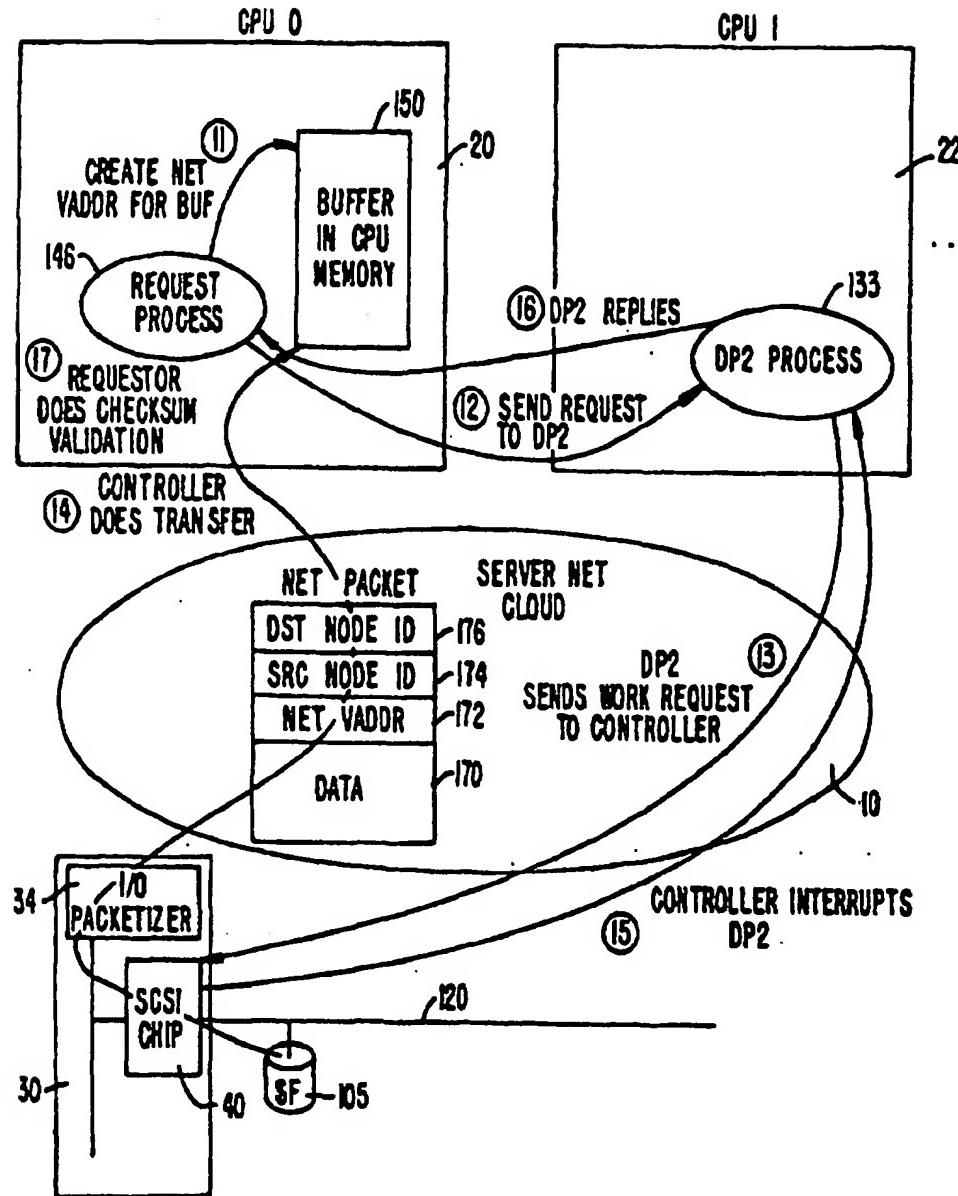


FIG. 3.



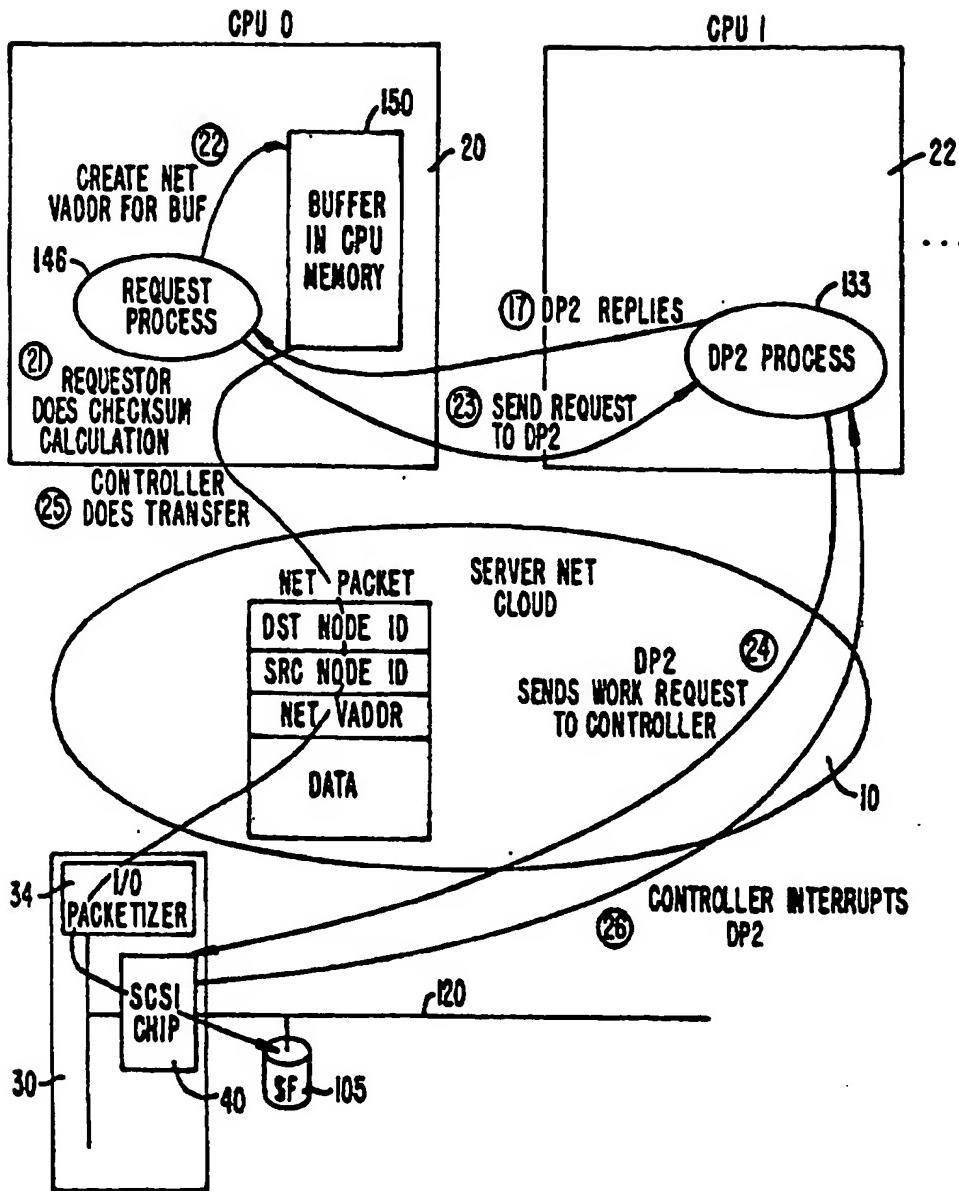


FIG. 5.

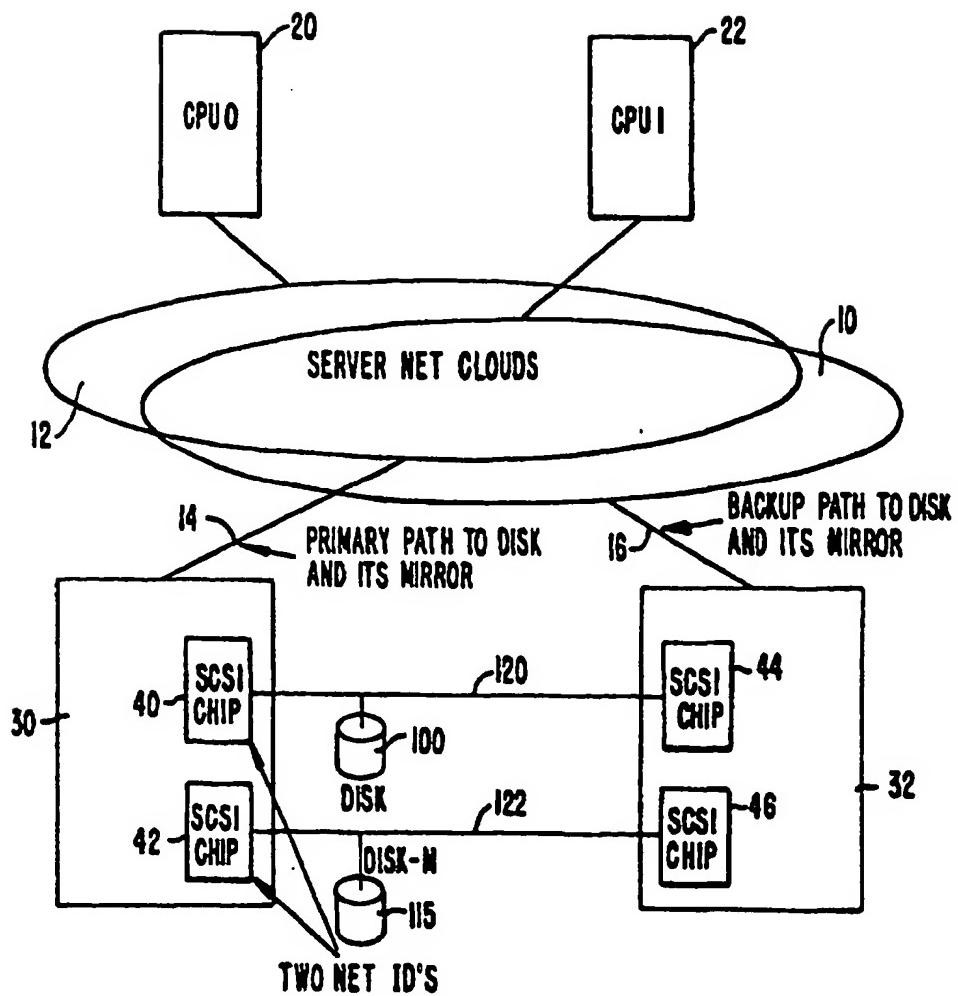


FIG. 6

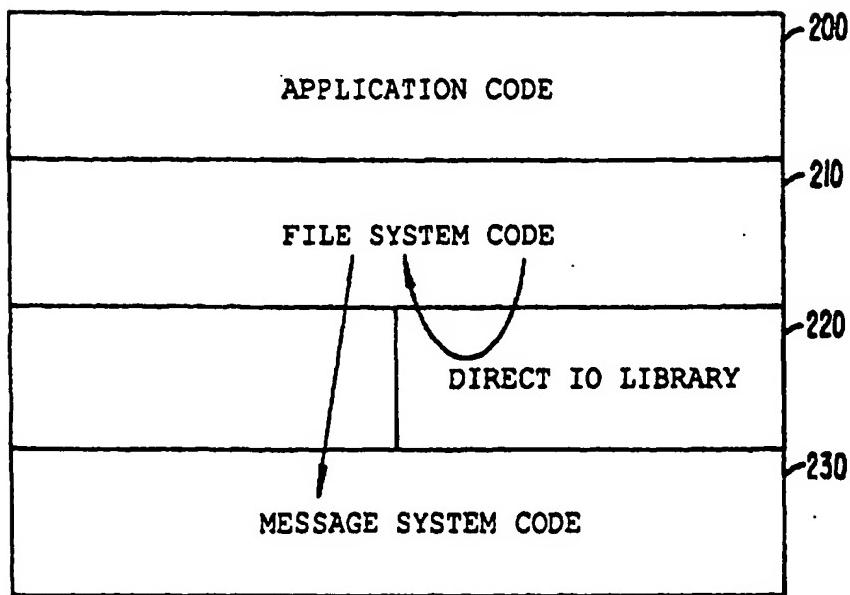


FIG. 7.

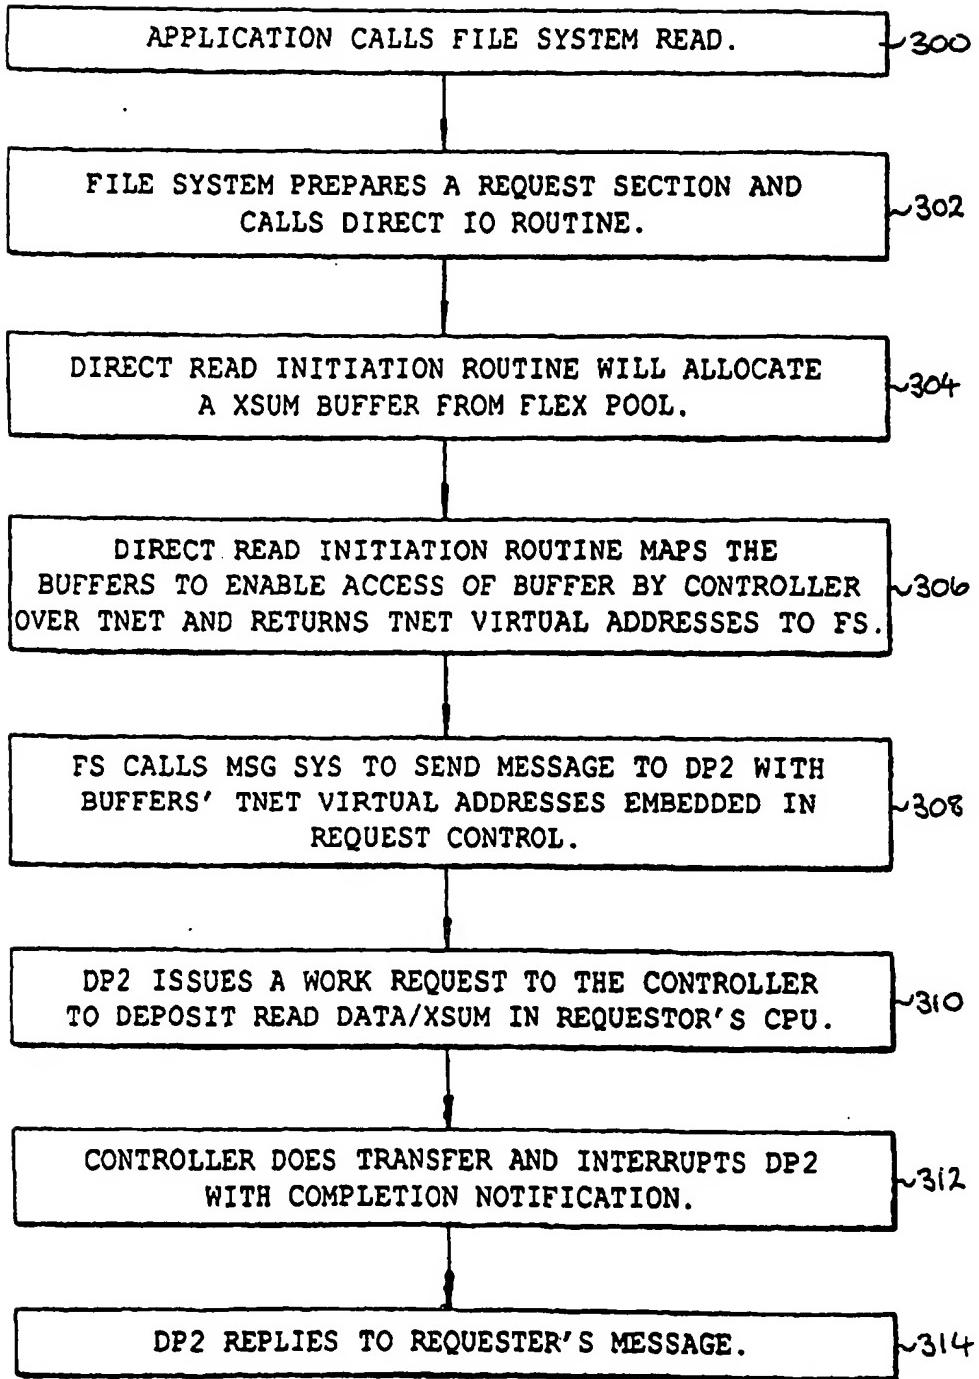


FIG. 8.

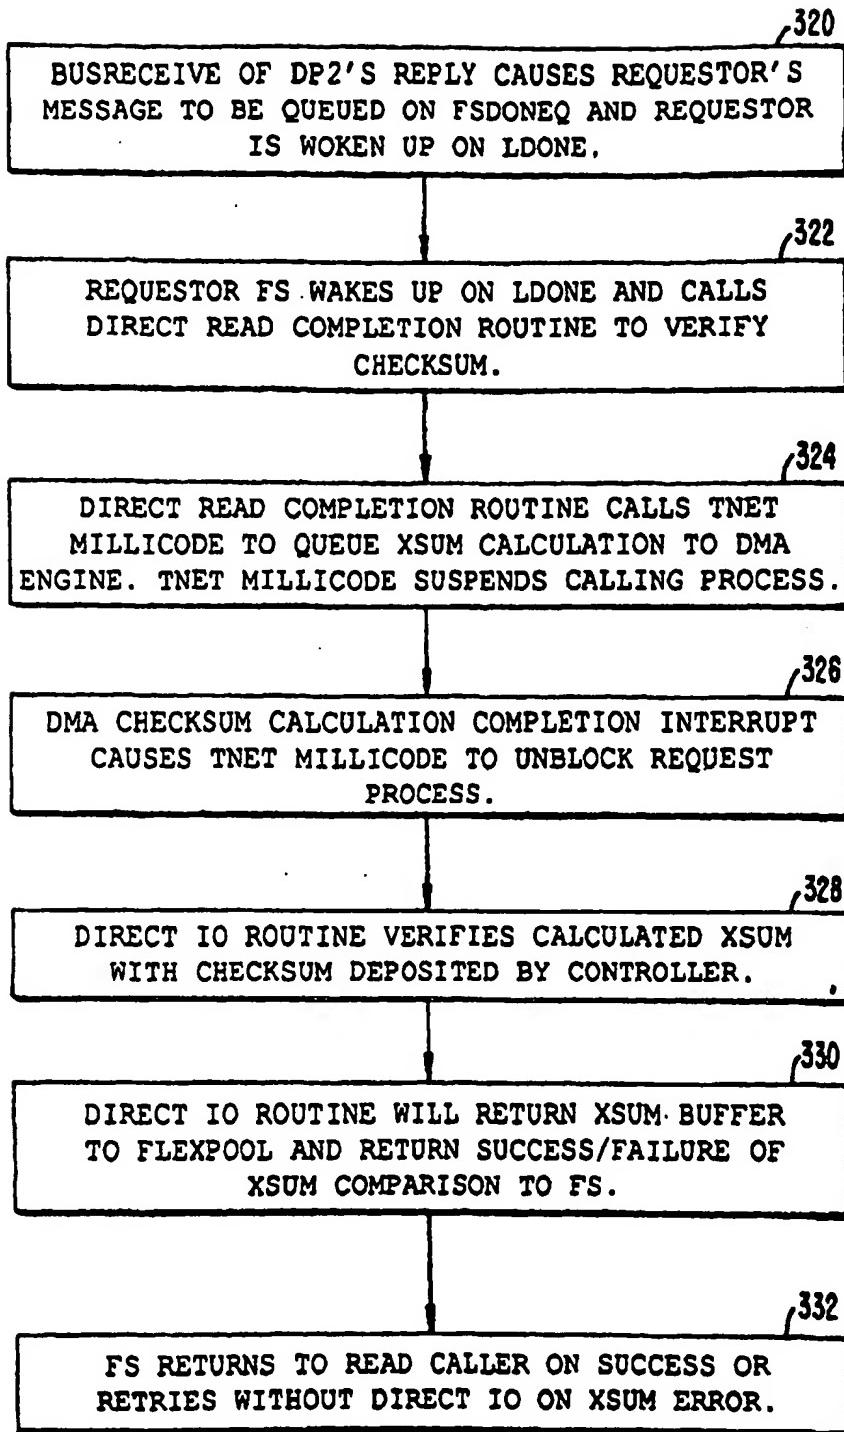


FIG. 9.

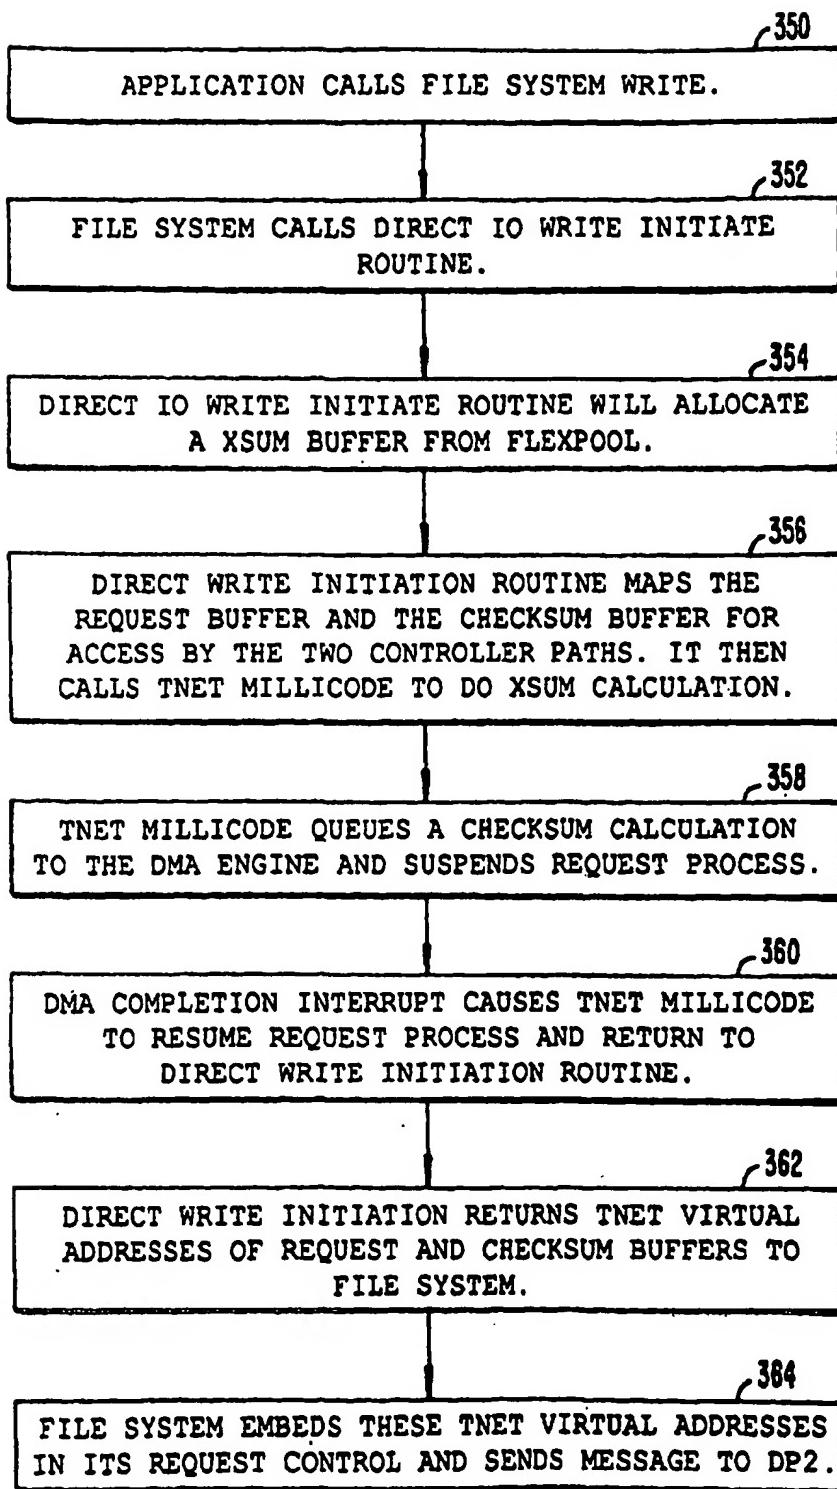


FIG. 10.

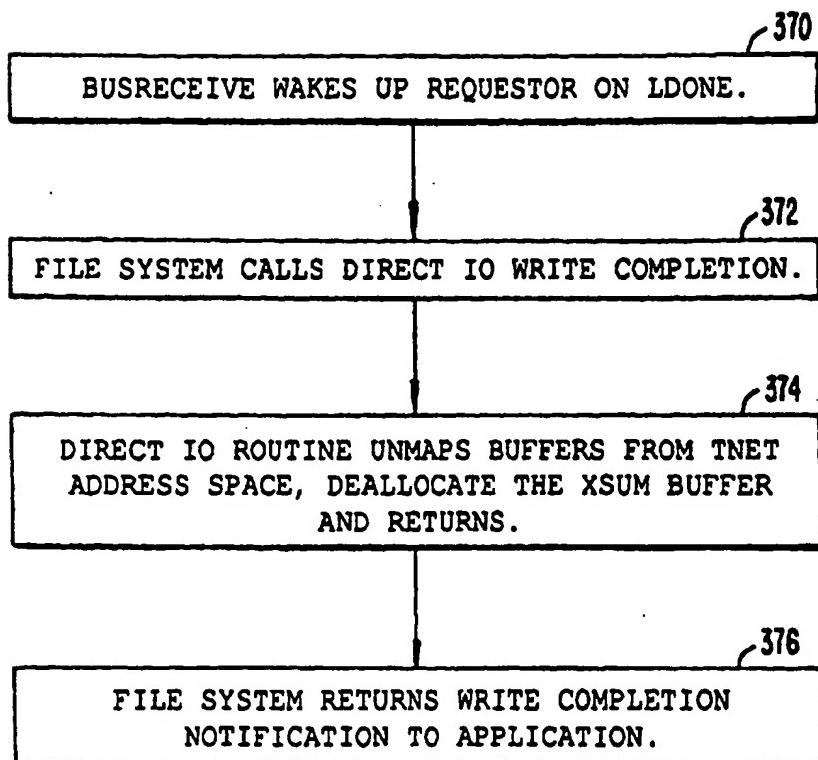


FIG. 11.